

Scalable Parallel Numerical CSP Solver

Daisuke Ishii¹, Kazuki Yoshizoe^{2,1} and Toyotaro Suzumura^{3,2}

¹ Tokyo Institute of Technology, Tokyo, Japan.

² Japan Science and Technology Agency, Japan.

³ IBM Research, Dublin, Ireland.

{dsksh,yoshizoe,suzumura}@acm.org

Abstract. We present a parallel solver for numerical constraint satisfaction problems (NCSPs) that can scale on a number of cores. Our proposed method runs worker solvers on the available cores and simultaneously the workers cooperate for the search space distribution and balancing. In the experiments, we attained up to 119-fold speedup using 256 cores of a parallel computer.

1 Introduction

Numerical constraint satisfaction problems (NCSPs, Section 2) have been successfully applied to problems described in the domain of reals[3,6]. Given a NCSP with search space represented as a *box* (i.e., interval vector), the *branch and prune algorithm* efficiently computes a *paving*, a set of boxes that encloses the solution set, yet its exponential computational complexity limits the tractable instances. Although the solving process exhibits a parallelism, no parallel NCSP solver has been made available to date because of the difficulty in partitioning the search space equally.

In this research, we parallelize a NCSP solver to scale its solving process on both shared memory and distributed memory parallel computers (Section 3). Our parallel method consists of parallel worker solvers that solve a portion of search space on CPU cores and interact with neighbor workers via message passing for dynamic load balancing. We also propose a preprocess that accelerates the initial search space distribution by sending sets of boxes via static routing between the workers. We have implemented the method by extending the Re-alpaver solver using the X10 language to realize a process-level parallelization over a number of cores. Section 4 reports experimental results when our method was deployed on two hardware environments.

Related work. There have been several work regarding parallel solving of CSPs with either discrete or continuous domains. Parallel solving of generic CSPs on massive computer clusters and supercomputers has been explored in [12,2,18]. This work focuses on a massive parallel solver for NCSPs that has a different characteristics compared to generic CSPs. In the survey [7], existing work is classified into (i) search-space splitting methods[16,12,2,5,15,18], (ii) cooperative methods for heterogeneous solver agents (cf. portfolios)[2], and (iii) parallelization of constraint propagation. Our method belongs to the first category. A few

works have used approaches (ii)[11] and (iii)[8] for parallelization of NCSP solving. However, to the best of our knowledge, a massive parallelization method that uses the typical approach (i) has not yet been proposed.

Substantial work regarding the parallelization of the *branch and bound* algorithm with search-space splitting exists [9,13,14]. This approach has also been applied to CSP solvers [16,5] and SAT solvers [15]. This work explores an efficient parallel method for solving NCSPs with similar approach to [9,13].

2 Numerical Constraint Satisfaction Problems

A *numerical constraint satisfaction problem* (NCSP) is defined as a triple (v, \mathbf{v}_0, c) that consists of a vector of *variables* $v = (v_1, \dots, v_n)$, an *initial domain* in the form of a box $\mathbf{v}_0 \in \mathbb{IR}^n$ (\mathbb{IR} denotes the set of closed real intervals), and a *constraint* $c(v) \equiv f(v) = 0 \wedge g(v) \geq 0$, where $f : \mathbb{R}^n \rightarrow \mathbb{R}^e$ and $g : \mathbb{R}^n \rightarrow \mathbb{R}^i$, i.e., a conjunction of e equations and i inequalities. A *solution* of a NCSP is an assignment of its variables $v \in \mathbf{v}_0$ that satisfies its constraints. The *solution set* Σ of a NCSP is the region within its initial domain that satisfies its constraints, i.e., $\Sigma(\mathbf{v}_0) := \{\tilde{v} \in \mathbf{v}_0 \mid c(\tilde{v})\}$. The target of this paper is *under-constrained* NCSPs such that $n > e$. In general, an under-constrained problem may have a continuous set of infinitely many solutions.

Branch and Prune Algorithm. The *branch and prune algorithm* [17] is the standard solving method for NCSPs. It takes a NCSP and a precision ϵ as an input and outputs a set of boxes (or *paving*) S that approximates the solution set with precision ϵ . Examples of S are illustrated in Figure 1.

An intermediate state of the algorithm is represented as a pair of sets of boxes (L, S) . The solver receives an initial state $(\{\mathbf{v}_0\}, \emptyset)$ and iteratively applies the *step computation* (illustrated in Figure 2) until it reaches a final state (\emptyset, S) . In the step computation, first, it takes the first element of the queue L of boxes and applies the **Prune** procedure, which is a filtering procedure that shaves boundary portions of the considered box. In this work, we use an implementation proposed in [6] which provides a verification process based on an interval Newton method combined with a relatively simple filtering process based on the Hull consistency[1]. As a result, a box becomes either empty, precise enough (its width is smaller than ϵ), verified as an *inner* box of the solution set Σ , or undecided. Precise and inner boxes are appended to S and undecided boxes are passed to **Branch**. Second, the **Branch** procedure bisects the box at the midpoint along a component corresponding to one of the variables and the sub-boxes will be put back in the queue. In this work, we assume **Branch** selects variables in an order which makes the search to behave in a breadth-first manner and thus the solving process gradually improves the precision of the overall pavings (Figure 1).

The computation of **Prune** is expensive and is the bottleneck of the solving process. Under certain conditions, application of **Prune** contracts a large portion of the search space into a tight box (cf. quadratic convergence of the interval Newton methods). **Prune** can also filter out the whole box if the considered box is

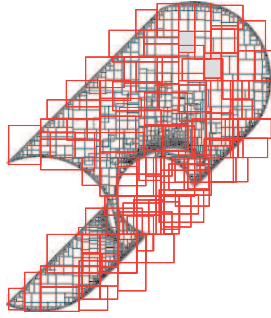


Fig. 1. Overlay of two solution box sets (pavings) with $\epsilon = 0.01, 0.1$

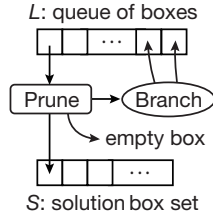


Fig. 2. A step computation of the branch and prune algorithm

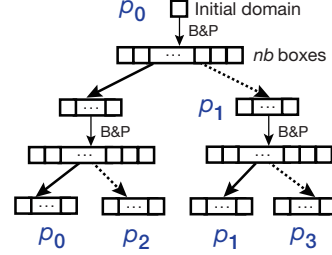


Fig. 3. Distribution of the box queue L in the preprocess

verified as an inner or totally inconsistent region. These characteristics of **Prune** result in the unbalanced nature of search trees. Therefore, a straightforward parallel method does not work efficiently. It is crucial for efficient NCSP solving to execute **Prune** on each step of traversing the search tree which makes it more difficult to distribute a search path among processors. These properties will be discussed in Section 4.1.

3 Parallel Branch and Prune

We propose a parallel method that runs several workers on the available CPU cores $p_0, \dots, p_{\#p-1}$ (we assume a single worker runs on a core and identify both a worker and a core with p_i). Workers homogeneously interleave the following three procedures and cooperate in a decentralized fashion: (i) breadth-first branch and prune search, (ii) distribution and workload balancing of search space in a sender-initiated manner, and (iii) termination detection. Sub-trees in the search space of the branch and prune becomes unbalanced but can be searched independently: there are no confluence of multiple branches, and we have no shared information between branches. Distribution of the search space among workers is done in preprocess and postprocess as described in the following subsections. The preprocess distributes the portions of the search space to the other workers, and the postprocess balances the load of each worker during search. Termination is detected by circulating a termination token via idling workers based on Dijkstra's method (see Section 11.4.4 of [9]).

Preprocess: Search-Space Splitting and Distribution. A solving process is started by a worker p_0 that possesses the initial domain (i.e., a box which contains the whole search space) in the queue L . To distribute the subsequent search space (i.e., a queue of boxes) equally to each core, the preprocess invokes a partition of the queue in two (or more) and then sends a portion to another worker. Figure 3 illustrates in a downward direction some initial transitions of the box queue L distributed among three workers. In our implementation, the

distribution routing is formed as a binary tree whose height is $\lceil \log_2 \#p \rceil$. In each node of the tree, the branch and prune process runs until the number of boxes in the queue reaches nb . Next, the queue is sorted by the volume of the boxes and the half of the content (i.e., $nb/2$ boxes) is sent to the other core via the right branch.

Postprocess: Dynamic Load Balancing. During search, each worker normalizes the loads within a predefined *neighborhood* which consists of a small number of *neighbor* workers. Because there are sufficiently large number of boxes, we simply regard the number of boxes in the queue L as the amount of load. Assume $\#p$ workers are running and each worker p_i possesses l_i boxes in its queue. We also assume for each worker p_i that N_i is a set of $|N|$ neighbor workers, N_i^{-1} is a set of workers where $p_j \in N_i^{-1} \Leftrightarrow p_i \in N_j$, L_i is a set of loads of the neighbor workers, and Δ is a predefined load margin. The load balancing procedure of a worker p_i performs the following steps once every ns branch and prune steps.

1. For each worker p_j in N_i^{-1} , inform the load l_i and put in the list L_j .
2. Calculate the mean μ of the loads in L_i .
3. If $\mu < \Delta$, for each worker p_j in N_i , send at most $\mu - l_j$ boxes to p_j (to be efficient, a certain number of boxes should be kept locally).

Neighbor workers can be identified e.g. as adjacent nodes in the $|N|$ -dimensional mesh of workers. The routing between neighbor workers is fixed during a solving process and thus it may happen that a worker possesses an excessive load than others. However, this load imbalance will be resolved by the subsequent load balancing processes.

4 Experimental Results

We have implemented the proposed method and measured the speedup of the solving process of under-constrained NCSPs. The experiments were performed with an exhaustive set of parameter combinations to explore the optimal settings.

Implementation. We have implemented the proposed method with C++ (gcc ver. 4.4.7 and 4.3.4) and X10 (ver. 2.3.1)[4], a high productivity language for parallel computing. In the following, we use the term *place*, which is a notion of X10 that in our setting represents a CPU core. Libraries Realpaver (ver. 1.1)[10] for sequential NCSP solving and Gaol (ver. 4.0.1)⁴ for basic interval computation were used to facilitate the implementation. The **Prune** procedure was realized by calling the sequential implementation in Realpaver. Each run of **Prune** took around 0.2–1ms and the overall execution became the bottleneck of the branch and prune algorithm (occupied greater than 95% of running time in sequential solving). The procedures for search space distribution and load balancing were implemented with X10. Communication of boxes and loads between places were

⁴ <http://sourceforge.net/projects/gaol/>

implemented as `async` tasks and performed in parallel to the search process so that the overhead will be hidden. In the experiments, timings t_1 for sequential runs on single core were measured using the C++ implementation described in [6], which worked identically and faster than our X10 version.

Experiment Environments. Two sets of experiments were operated using (1) a shared-memory machine equipped with 40 cores (four of 10-core Intel Xeon E7-4860 2.26GHz) and 256GB of local memory and (2) up to 256 cores of SGI UV1000, a pseudo-shared memory machine equipped with 2,048 cores (8-core Xeon E7-8837 2.67GHz) and 16TB of memory. UV1000 works as a single shared memory machine by emulating memory accesses using communication based on a high speed NUMalink5 network which has a bandwidth of 120Gbps. We used the MPI backend of X10 with options `X10_NTHREADS = 6` and `GC_NPROCS = 2`.

Experiments on a Shared Memory Machine. We solved the problems shown in [6,3] using 40 cores of the machine (1). We report the results for two representative problems. Parameters in the load balancing method were set as either combination of the following values: $nb = 32$, $|N| \in \{2, 4\}$, $ns \in \{10, 100, 1000\}$, and $\Delta = 10$. We also computed with and without the preprocess (when the preprocess is not used, the postprocess is executed from the beginning). For each problem, we solved two instances with two multiplicative precisions. The specification of each instance and the computational results are presented in Table 1. In the table, the columns “problem”, “size”, “ ϵ ”, “pp”, and “ $|N|$ ” represent the name of the problem, size (i.e., the number of projection/parameter variables), the precision, usage of the preprocess, and the number of neighbor workers, respectively. The rest of the columns represents the results. t_1 and $\#br_1$ represent the running time and the number of branches on single core. $t_{ns,i}$ and $\#br_{ns,i}$ represent the running time and the largest number of branches performed by a worker when computed with the interval ns and i X10 places (best timings are underlined). Figure 4 illustrates the speedup of the solving process.

Experiments on a Cluster with High-speed Interconnection. We solved the problem “*3rpr*” using up to 256 cores of the machine (2), UV1000. Parameters in the load balancing method were set as either combination of the following values: $nb = 8$, $|N| \in \{2, 4\}$, $ns = 1000$, and $\Delta = 10$. The results are presented in Table 2. Each column of the table represents the same information as presented in Table 1 except that the column “ $\#sends_{1000,256}$ ” represents the number of loads sent by the load balancer in the solving process with $ns = 1000$ and 256 X10 places. Figure 5 illustrates the speedup of the solving process.

4.1 Discussions

In the experiments, our method scaled up to 256 cores with the optimal configurations. We achieved speedups up to 32.3 fold using 40 cores of the shared memory machine and up to 119 fold using 256 cores of the cluster machine.

The best speedup of 119 fold was obtained with the preprocess. The preprocess facilitates and accelerates the workload distribution in the early stage

Table 1. Experimental result on the shared memory machine

problem	size	ϵ	pp	$ N $	t_1	$t_{10,40}$	$t_{100,40}$	$t_{1000,40}$	#br ₁	$\frac{\#br_1}{\#br_{100,40}}$
4D sphere and plane (<i>sp2-4</i>)	2+4	0.02	yes	2	254	9.52	<u>7.94</u>	10.1	238 319	37.0
				4		11.0	8.73	9.45		37.6
				2		9.34	8.13	16.5		35.4
				4		10.8	8.67	10.5		37.3
		0.01	yes	2	721	38.6	22.8	23.9	669 601	38.0
				4		38.1	25.1	24.0		38.7
				2		35.7	<u>22.7</u>	30.0		37.6
				4		36.8	24.7	24.7		38.7
3-RPR robot (<i>3rpr</i>)	3+3	0.2	yes	2	1 100	199	73.9	34.1	1 936 939	33.7
				4		296	68.5	36.4		38.0
				2		185	64.0	<u>34.0</u>		33.5
				4		244	58.3	36.1		38.0
		0.1	yes	2	4 080	1 010	714	282	7 186 845	30.2
				4		2 820	1 070	257		36.0
				2		971	678	244		28.5
				4		2 630	901	<u>231</u>		36.0

Table 2. Experimental result on the UV1000 cluster

problem	size	ϵ	pp	$ N $	t_1	$t_{1000,32}$	$t_{1000,256}$	$\frac{\#br_1}{\#br_{1000,256}}$	#sends _{1000,256}
3-RPR robot (<i>3rpr</i>)	3+3	0.2	yes	2	850	37.4	<u>14.0</u>	131	18 648
				4		54.6	33.6	157	85 904
				2		39.8	20.4	43.8	10 856
				4		53.4	32.0	123	66 212
		0.1	yes	2	3 040	341	<u>25.6</u>	192	39 608
				4		371	87.8	176	128 084
				2		325	54.7	105	34 892
				4		339	52.1	184	129 512

of the search process. In some of the experiments without using the preprocess, the speedup ratio became saturated when using many cores (e.g., *sp2-4* with $ns = 1000$ and the experiments on the cluster). This was because the load balancing process was too infrequent for the given number of workers and the work load diffusion became too slow. When comparing the right-hand graphs for the instance *sp2-4*, $ns = 1000$, in Figure 4, we can notice that the point of saturation shifts according to the search space size. On the other hand, in some other experiments, the results got worse with the preprocess (e.g., results with $ns = 10$ on the machine (1)). It occasionally happens that the preprocess mostly solves the problem. However, the preprocess can result in highly unbalanced search trees because of the Prune process, and in such cases the postprocess will not have enough time for load balancing.

Regarding the neighborhood sizes $|N| = 2, 4$, there was a trade off between the workloads balance and the amount of communications required. For the

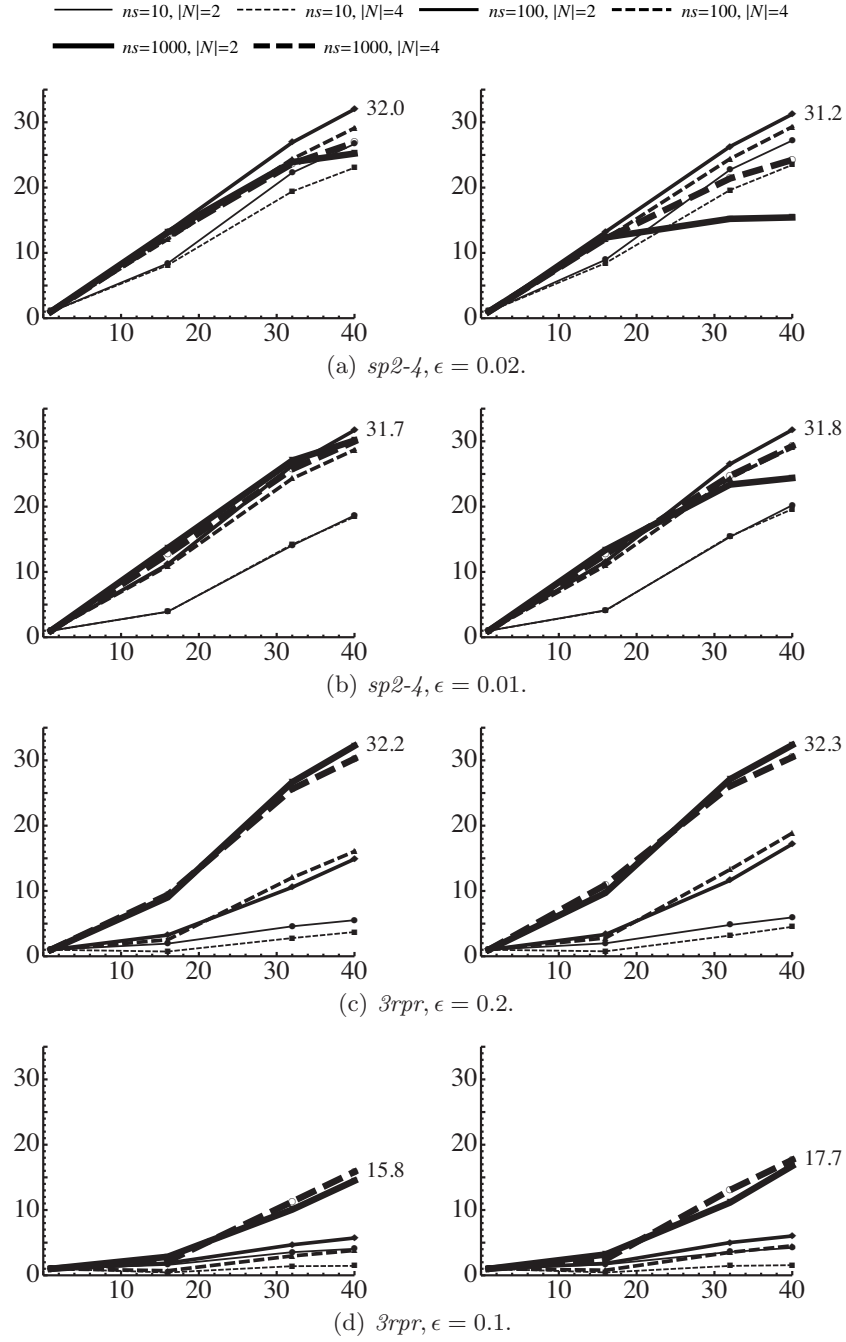


Fig. 4. Speedups on the shared memory machine with 40 cores. Left- and right-hand side graphs correspond to computations with and without the preprocess, respectively

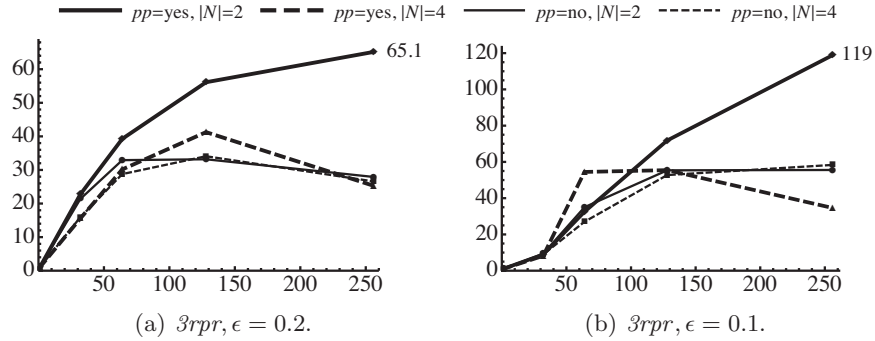


Fig. 5. Speedups using 256 cores of UV1000

shared memory machine, it was unclear which size had the advantage. However, for UV1000, the solver was notably slower for $|N| = 4$ than $|N| = 2$. It is understandable because larger number of neighbors significantly increased the number of communications (see “#sends_{1000,256}” in Table 2) and communications between places were much more costly compared to normal shared memory machines despite the high speed network of UV1000.

Three intervals $ns = 10, 100, 1000$ were used for load balancing which determined the speed of workload distribution. When the distribution was too slow, the speedup ratio did not scale well (e.g., $3rpr$, $\epsilon = 0.2$, with $ns = 1000$ on the machine (1)). Conversely, small intervals required greater amount of communications and therefore we used $ns = 1000$ to draw better performance on the cluster where communications were more costly.

There was a large overhead caused by the workers sending a large number of boxes for load balancing when the number of workers was not sufficient against the problem size. Speedups for $3rpr$, $\epsilon = 0.1$, using 40 workers or less shows an example of such overheads (Figure 4(d)). Resolving this overhead by suppressing redundant box sends is a part of the future work.

5 Conclusions

In this paper, we proposed a parallel branch and prune algorithm, based on, non-portfolio, search-space splitting approach. In the experiments, using 256 X10 places (i.e., cores), we achieved speedup factors of as much as 119. We expect that our parallelized solver will be applied to large practical problems, e.g., the robotics problems in [3].

Acknowledgments. This work was partially funded by JSPS (KAKENHI 25880008 and 25700038). Computing resources for the experiments in this paper were provided by Prof. Kazunori Ueda (Waseda University, Tokyo) and a Compute Canada RAC award, for environments (1) and (2), respectively.

References

1. Benhamou, F., Goualard, F., Granvilliers, L., Puget, J.F.: Revising Hull and Box Consistency. In: Proc. of ICLP. pp. 230–244 (1999)
2. Bordeaux, L., Hamadi, Y., Samulowitz, H.: Experiments with Massively Parallel Constraint Solving. In: Proc. of IJCAI. pp. 443–448 (2006)
3. Caro, S., Chablat, D., Goldsztejn, A., Ishii, D., Jermann, C.: A branch and prune algorithm for the computation of generalized aspects of parallel robots. *Artificial Intelligence* 211, 34–50 (2014)
4. Charles, P., Grothoff, C., Saraswat, V.: X10: an object-oriented approach to non-uniform cluster computing. In: Proc. of OOPSLA. pp. 519–538 (2005)
5. Chu, G., Schulte, C., Stuckey, P.: Confidence-based work stealing in parallel constraint programming. In: Proc. of CP. pp. 226–241. *LNCIS* 5732 (2009)
6. Ishii, D., Goldsztejn, A., Jermann, C.: Interval-based projection method for under-constrained numerical systems. *Constraints Journal* 17(4), 432–460 (2012)
7. Gent, I.P., Jefferson, C., Miguel, I., Moore, N.C.A., Nightingale, P., Prosser, P., Unsworth, C.: A Preliminary Review of Literature on Parallel Constraint Solving. In: Proc. of Workshop on Parallel Methods for Constraint Solving. 13 pp. (2011)
8. Goldsztejn, A., Goualard, F.: Box consistency through adaptive shaving. Proc. of SAC. pp. 2049–2054 (2010)
9. Grama, A., Gupta, A., Karypis, G., Kumar, V.: Introduction to Parallel Computing. Addison Wesley (2003)
10. Granvilliers, L., Benhamou, F.: Algorithm 852: RealPaver: an interval solver using constraint satisfaction techniques. *ACM Transactions on Mathematical Software* 32(1), 138–156 (2006)
11. Granvilliers, L., Hains, G.: A conservative scheme for parallel interval narrowing. *Information Processing Letters* 74(3-4), 141–146 (2000)
12. Jaffar, J., Santosa, A., Yap, R., Zhu, K.: Scalable distributed depth-first search with greedy work stealing. In: Proc. of ICTAI. pp. 98–103 (2004)
13. Lüling, R., Monien, B., Reinefeld, A., Tschöke, S.: Mapping Tree-Structured Combinatorial Optimization Problems onto Parallel Computers. In: Proc. of Solving Combinatorial Optimization Problems in Parallel. pp. 115–144. *LNCIS* 7141 (1996)
14. Otten, L., Dechter, R.: Towards Parallel Search for Optimization in Graphical Models. In: Proc. of ISAIM. 8 pp. (2010)
15. Schubert, T., Lewis, M., Becker, B.: PaMiraXT: Parallel SAT Solving with Threads and Message Passing. *JSAT* 6, 203–222 (2009)
16. Schulte, C.: Parallel search made simple. In: Proc. of TRICS. pp. 41–57 (2000)
17. Van Hentenryck, P., McAllester, D., Kapur, D.: Solving Polynomial Systems Using a Branch and Prune Approach. *SIAM Journal on Numerical Analysis* 34(2), 797–827 (1997)
18. Xie, F., Davenport, A.: Massively Parallel Constraint Programming for Supercomputers : Challenges and Initial Results. In: Proc. of CPAIOR. pp. 334–338. *LNCIS* 6140 (2010)